Novelty Assessment Report

Paper: The Coverage Principle: How Pre-Training Enables Post-Training
PDF URL: https://openreview.net/pdf?id=AUXvYQlQLZ
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-27

Abstract

Language models demonstrate remarkable abilities when pre-trained on large text corpora and fine-tuned for specific tasks, but how and why pre-training shapes the success of the final model remains poorly understood. Notably, although pre-training success is often quantified by cross entropy loss, cross entropy can be poorly predictive of downstream performance. Instead, we provide a theoretical perspective on this relationship through the lens of coverage, which quantifies the probability mass the pre-trained model places on high-quality responses and which is necessary and sufficient for post-training and test-time scaling methods like Best-of-N to succeed. Our main results develop an understanding of the coverage principle, a phenomenon whereby next-token prediction implicitly optimizes toward a model with good coverage. In particular, we uncover a mechanism that explains the power of coverage in predicting downstream performance: coverage generalizes faster than cross entropy, avoiding spurious dependence on problem dependent parameters such as the sequence length. We also study practical algorithmic interventions with provable benefits for improving coverage, including (i) model/checkpoint selection procedures, (ii) gradient normalization schemes, and (iii) test-time decoding strategies.

Core Task Landscape

This paper addresses: **Understanding How Pre-Training Enables Post-Training Success Through Coverage Optimization**

A total of **35 papers** were analyzed and organized into a taxonomy with **32 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Coverage and Generalization Theory**
- **Training Methodologies and Optimization**
- **Post-Training Alignment and Preference Optimization**
- **Application Domains and Task-Specific Adaptation**
- **Model Compression and Efficiency**
- **Catastrophic Forgetting and Knowledge Retention**
- **Zero-Shot and Transfer Learning**
- **Survey and Review Literature**
- **Specialized Post-Training Techniques**
- **Downstream Task Optimization**
- ... and 1 more categories

Complete Taxonomy Tree

- Understanding How Pre-Training Enables Post-Training Success Through Coverage Optimization Survey Taxonomy
- Coverage and Generalization Theory
  - Coverage Principle and Mechanisms ★ (2 papers)
  - [0] The Coverage Principle: How Pre-Training Enables Post-Training (Anon et al., 2026) View paper
  - [11] Self-improvement in language models: The sharpening mechanism (Huang Audrey, 2024) View paper
  - Architectural Generalization (1 papers)
  - [18] Understanding and Improving Length Generalization in Recurrent Models (Gu, 2025) View paper
- Training Methodologies and Optimization
  - Data-Efficient Training Paradigms
  - Minimal-Data Reasoning (1 papers)
    - [1] Limo: Less is more for reasoning (Ye Yixin, 2025) View paper
  - Small-Scale Reasoning Models (1 papers)
    - [9] Mobilellm-r1: Exploring the limits of sub-billion language model reasoners with open training recipes (Zhao Changsheng, 2025) View paper
  - Multi-Stage Post-Training Frameworks
  - Agentic and Reinforcement Learning Pipelines (1 papers)
    - [2] Kimi k2: Open agentic intelligence (Kimi Team, 2025) View paper
  - Self-Improving Systems (1 papers)
    - [3] Will pre-training ever end? a first step toward next-generation foundation mllms via self-improving systematic cognition (Zhang Xiao-ying, 2025) View paper
  - Training Dynamics and Stage Analysis
  - Comprehensive Stage Evaluation (1 papers)
    - [25] EvoLM: In Search of Lost Language Model Training Dynamics (Qi, 2025) View paper

    - Controlled Causal Analysis (1 papers)
        - [26] On the Interplay of Pre-Training, Mid-Training, and RL on Reasoning Language Models (Charlie Zhang, 2025) View paper
    - Domain Adaptation and Specialization
    - Financial Domain Adaptation (1 papers)
        - [17] Demystifying Domain-adaptive Post-training for Financial LLMs (Ke, 2025) View paper
    - Mathematical Reasoning Enhancement (1 papers)
        - [10] Advancing Mathematical Reasoning in Language Models: The Impact of Problem-Solving Data, Data Synthesis Methods, and Training Stages (Chen, 2025) View paper
- Post-Training Alignment and Preference Optimization
    - Preference-Based Alignment (1 papers)
    - [4] Direct Post-Training Preference Alignment for Multi-Agent Motion Generation Models Using Implicit Feedback from Pre-training Demonstrations (Tian Ran, 2025) View paper
    - Robustness-Aware Alignment (1 papers)
    - [5] RobustVLA: Robustness-Aware Reinforcement Post-Training for Vision-Language-Action Models (Zhang HongYin, 2025) View paper
    - Distributional RL for Post-Training (1 papers)
    - [28] : Provably Optimal Distributional RL for LLM Post-Training (JP Zhou, 2025) View paper
    - Inverse Reward Learning (1 papers)
    - [31] Reinforcement Learning with Inverse Rewards for World Model Post-training (Ye Yang, 2025) View paper
- Application Domains and Task-Specific Adaptation
    - Embodied AI and Robotics (1 papers)
    - [20] Poutine: Vision-Language-Trajectory Pre-Training and Reinforcement Learning Post-Training Enable Robust End-to-End Autonomous Driving (Rowe, 2025) View paper
    - Dialogue Systems (1 papers)
    - [21] Improving Retrieval-Based Dialogue Systems: Fine-Grained Post-training Prompt Adaptation and Pairwise Optimization Fine-Tuning Strategy (Tianqing Zhang, 2024) View paper
    - Precipitation Forecasting (1 papers)
    - [32] Self-supervised Pre-training for Precipitation Post-processor (Sojung An, 2023) View paper
    - Enterprise Decision Support (1 papers)
    - [24] Intelligent optimization of transfer learning and knowledge tracking in enterprise strategic decision support systems (Wensheng Yan, 2025) View paper
- Model Compression and Efficiency
    - Post-Training Quantization (1 papers)
    - [6] GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers (Frantar, 2022) View paper
    - Feature Quantization (1 papers)
    - [30] Improving the Post-Training Neural Network Quantization by Prepositive Feature Quantization (Tianshu Chu, 2023) View paper
    - RNS-Based Sparsity (1 papers)
    - [34] Improving Residue-Level Sparsity in RNS-based Neural Network Hardware Accelerators via Regularization (E. Kavvousanos, 2023) View paper
- Catastrophic Forgetting and Knowledge Retention
    - Layer Scaling for Forgetting Prevention (1 papers)
    - [12] LiNeS: Post-training Layer Scaling Prevents Forgetting and Enhances Model Merging (Wang Ke, 2024) View paper
    - Pre-Training Poisoning Persistence (1 papers)
    - [8] Persistent Pre-Training Poisoning of LLMs (Zhang Yiming, 2024) View paper
    - Knowledge Extraction and Parameterization (1 papers)
    - [27] Information Extraction-Driven Knowledge-Based Salient Masked Pre-Training Method (Shaochong Lei, 2025) View paper
- Zero-Shot and Transfer Learning
    - Image-Free Zero-Shot Classification (1 papers)
    - [14] Image-free Classifier Injection for Zero-Shot Classification (Anders Christensen, 2023) View paper
    - Post-Deployment Adaptation (1 papers)
    - [15] Post-Deployment Adaptation with Access to Source Data via Federated Learning and Source-Target Remote Gradient Alignment (Wagner, 2023) View paper
- Survey and Review Literature (3 papers)
    - [7] Large Language Models Post-training: Surveying Techniques from Alignment to Reasoning (G Tie, 2025) View paper
    - [19] A Survey of Post-Training Scaling in Large Language Models (Hanyu Lai, 2025) View paper
- Specialized Post-Training Techniques
    - BERT and Transformer Post-Training (1 papers)
    - [35] A Robustly Optimized BERT Pre-training Approach with Post-training (Zhuang Liu, 2021) View paper
    - Last-Layer Training (1 papers)
    - [33] Post Training in Deep Learning with Last Kernel (Moreau, 2022) View paper
    - Voltage Trajectory Prediction (1 papers)
    - [22] Conformalized prediction of post-fault voltage trajectories using pre-trained and finetuned attention-driven neural operators. (Amirhossein Mollaali, 2025) View paper
- Downstream Task Optimization (1 papers)
    - [29] Optimising Language Models for Downstream Tasks: A Post-Training Perspective (Shi Zheng-yan, 2025) View paper
- Human Factors and Training Effectiveness (2 papers)
    - [13] Predicting post-training reactions from pre-training attitudes (Adomaityte, 2025) View paper
    - [23] A field study of computer efficacy beliefs as an outcome of training: the role of computer playfulness, computer knowledge, and performance during training (Denise Potosky, 2002) View paper

Narrative

Core task: understanding how pre-training enables post-training success through coverage optimization. The field has evolved into a rich landscape organized around several complementary perspectives. Coverage and Generalization Theory examines foundational principles governing how pre-trained representations support downstream adaptation, while Training Methodologies and Optimization explores algorithmic strategies for effective learning across stages. Post-Training Alignment and Preference Optimization focuses on steering models toward desired behaviors, often through human feedback or reward signals. Application Domains and Task-Specific Adaptation investigates how these principles manifest in specialized settings such as finance, vision-language tasks, and mathematical reasoning. Meanwhile, Model Compression and Efficiency addresses resource constraints, Catastrophic Forgetting and Knowledge Retention studies stability during continual learning, and Zero-Shot and Transfer Learning probes generalization without explicit fine-tuning. Survey and Review Literature synthesizes emerging insights, while Specialized Post-Training Techniques and Downstream Task Optimization refine methods for particular scenarios.

Within this ecosystem, several active lines of work reveal key trade-offs. Many studies explore how pre-training coverage shapes post-training efficiency, examining whether broad exposure during pre-training reduces the need for extensive downstream data or whether targeted mid-training interventions can bridge gaps. Others investigate alignment mechanisms that balance capability retention with preference learning, as seen in works like Direct Preference Alignment[4] and Self-Improving Systematic Cognition[3]. The Coverage Principle[0] sits squarely within the Coverage and Generalization Theory branch, offering a mechanistic lens on how pre-training diversity enables robust post-training outcomes. Its emphasis on coverage optimization complements neighboring work such as Sharpening Mechanism[11], which examines how post-training refines pre-trained features. Together, these contributions illuminate the interplay between pre-training breadth and post-training specialization, a central question as the field seeks to understand what makes certain pre-trained models more amenable to downstream success.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Self-improvement in language models: The sharpening mechanism

Authors: Huang Audrey, Block, Adam, Audrey Huang, Foster, et al. (23 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Recent work in language modeling has raised the possibility of self-improvement, where a language models evaluates and refines its own generations to achieve higher performance without external feedback. It is impossible for this self-improvement to create information that is not already in the model, so why should we expect that this will lead to improved capabilities? We offer a new perspective on the capabilities of self-improvement through a lens we refer to as sharpening. Motivated by the o...

Relationship Analysis

Both papers belong to the Coverage Principle and Mechanisms category, investigating how coverage relates to downstream performance in language models. They share overlapping focus on coverage as a predictor of post-training success and its role in enabling test-time scaling methods like Best-of-N sampling. However, the original paper provides a theoretical framework explaining why next-token prediction implicitly optimizes coverage and develops generalization bounds, while the candidate paper focuses on self-improvement through a "sharpening mechanism" that uses the model as its own verifier to amortize inference-time computation, analyzing specific algorithms like SFT-Sharpening and RLHF-Sharpening.

Contributions Analysis

Overall novelty summary. The paper proposes that coverage—the probability mass a pre-trained model assigns to high-quality responses—is a better predictor of downstream success than cross-entropy loss. It sits within the Coverage Principle and Mechanisms leaf, which contains only one sibling paper examining sharpening mechanisms in post-training. This leaf is part of the broader Coverage and Generalization Theory branch, which itself comprises just two leaves with three total papers. The sparse population suggests this theoretical perspective on pre-training success is relatively underexplored compared to the more crowded Training Methodologies branch containing over twenty papers across multiple subtopics.

The taxonomy reveals substantial activity in neighboring areas. The Training Methodologies and Optimization branch encompasses data-efficient paradigms, multi-stage frameworks, and domain adaptation, with papers examining how pre-training interacts with supervised fine-tuning and reinforcement learning. Post-Training Alignment explores preference optimization and robustness-aware methods. The paper's focus on coverage as a unifying principle bridges these areas: it provides theoretical grounding for why certain pre-training strategies enable effective post-training, whether through Best-of-N sampling or alignment techniques. However, the taxonomy shows limited prior work explicitly connecting coverage theory to these downstream applications.

Among twenty-four candidates examined, none clearly refute the three main contributions. The coverage principle contribution examined six candidates with zero refutations. The generalization analysis showing coverage generalizes faster than cross-entropy examined nine candidates, again with no refutations. The algorithmic interventions contribution also examined nine candidates without finding overlapping prior work. These statistics suggest the theoretical framing through coverage is relatively novel within the limited search scope, though the small candidate pool means potentially relevant work in adjacent areas—such as scaling laws or representation learning—may not have been captured.

The analysis reflects a focused but limited literature search rather than exhaustive coverage of pre-training theory. The sparse taxonomy structure around coverage principles and the absence of refuting candidates among twenty-four examined papers suggest the work occupies a relatively unexplored theoretical niche. However, the broader Training Methodologies branch shows substantial empirical work on pre-training and post-training interactions, indicating the practical phenomena this paper theorizes about are well-studied, even if the coverage-centric lens is less common.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: The coverage principle for next-token prediction

Description: The authors introduce the coverage profile as a novel metric that refines cross-entropy and show that next-token prediction implicitly optimizes toward models with good coverage. They prove that coverage generalizes faster than cross entropy, avoiding spurious dependence on problem-dependent parameters like sequence length.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. TrajTok: Technical Report for 2025 Waymo Open Sim Agents Challenge

URL: View paper

Brief Assessment

TrajTok[48] focuses on trajectory tokenization for autonomous driving simulation, not on theoretical analysis of next-token prediction or coverage metrics for language models. The paper addresses discrete token generation for vehicle trajectories rather than the statistical properties of pre-training objectives.

---

2. Human-Inspired Learning for Large Language Models via Obvious Record and Maximum-Entropy Method Discovery

URL: View paper

Brief Assessment

Human-Inspired Learning[49] focuses on explicit symbolic memory and entropy-guided method discovery for rare scenarios, not on coverage metrics for next-token prediction or generalization analysis of pre-training objectives.

3. On the Generalization Ability of Next-Token-Prediction Pretraining

URL: View paper

Brief Assessment

Next-Token Generalization[51] focuses on generalization bounds for next-token prediction using Rademacher complexity and covering numbers, analyzing model parameters' effects. The original paper introduces coverage profile as a metric for downstream task success and proves coverage generalizes faster than cross-entropy. These are distinct technical contributions addressing different aspects of next-token prediction.

4. CoVeR: Conformal Calibration for Versatile and Reliable Autoregressive Next-Token Prediction

URL: View paper

Brief Assessment

CoVeR[50] focuses on conformal prediction for autoregressive decoding with token-level calibration and clustering, not on analyzing how pre-training with next-token prediction implicitly optimizes coverage profiles or their generalization properties relative to cross-entropy.

5. TFDP: Token-Efficient Disparity Audits for Autoregressive LLMs via Single-Token Masked Evaluation

URL: View paper

Brief Assessment

TFDP[46] focuses on single-token disparity audits for autoregressive LLMs through masked evaluation, not on coverage metrics for next-token prediction or generalization properties of language model training. The candidate addresses a completely different problem domain (bias detection and misinformation auditing) rather than pre-training optimization principles.

6. Language modeling techniques for biological sequence processing

URL: View paper

Brief Assessment

Biological Sequence Modeling[47] focuses on biological sequence processing using language modeling techniques for domain-specific applications (DNA, RNA, proteins). It does not address coverage metrics, next-token prediction theory, or the relationship between pre-training and post-training in language models, which are the core theoretical contributions of the original paper.

Contribution 2: Generalization analysis showing coverage generalizes faster than cross-entropy

Description: The authors develop a theoretical analysis (Theorem 4.1) demonstrating that maximum likelihood estimation achieves better generalization for coverage compared to cross-entropy, with rates that avoid dependence on sequence length and converge faster as the tail parameter N increases.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. On the sample complexity of next-token prediction

URL: View paper

Brief Assessment

Next-Token Sample Complexity[54] focuses on generalization bounds for next-token prediction in Markov chains using cross-entropy loss, not on comparing coverage versus cross-entropy generalization rates or maximum likelihood estimation properties as claimed in the original contribution.

2. A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses

URL: View paper

Brief Assessment

Mutual Information Metric[53] focuses on metric learning and establishing connections between cross-entropy and pairwise losses through mutual information theory. The candidate does not address generalization bounds for maximum likelihood estimation in sequence modeling contexts or coverage profiles as defined in the original paper.

3. Beyond maximum-likelihood training: analysis and methods for building robust language generation models

URL: View paper

Brief Assessment

Beyond Maximum-Likelihood[56] focuses on robustness issues in language generation (degeneration, repetition, toxicity) through imitation learning and entropy-centric analyses, not on generalization bounds comparing coverage versus cross-entropy convergence rates.

4. On how to avoid exacerbating spurious correlations when models are overparameterized

URL: View paper

Brief Assessment

Avoiding Spurious Correlations[61] focuses on worst-group error in imbalanced classification with spurious correlations, not on coverage profiles or maximum likelihood estimation generalization bounds. The candidate's theoretical framework addresses fairness in overparameterized models rather than pre-training/post-training dynamics.

5. Limits of sensing temporal concentration changes by single cells

URL: View paper

Brief Assessment

Temporal Concentration Sensing[62] focuses on biological concentration sensing accuracy and compares linear regression versus maximum likelihood estimation for detecting concentration ramps in chemotaxis. This is fundamentally different from analyzing generalization bounds for language model pre-training objectives.

6. Maximizing entropy on adversarial examples can improve generalization

URL: View paper

Brief Assessment

Maximizing Entropy Adversarial[60] focuses on improving generalization through entropy maximization on adversarial examples for classification tasks, not on analyzing maximum likelihood estimation or coverage metrics for language models.

7. Moment distributionally robust tree structured prediction

URL: View paper

Brief Assessment

Distributionally Robust Tree[59] focuses on tree-structured prediction with moment-based distributionally robust optimization and derives generalization bounds for worst-case risks. The original paper analyzes maximum likelihood estimation for language models with coverage profiles, a fundamentally different problem setting and theoretical framework.

8. Implicit Bias of Gradient Descent for Wide Two-layer Neural Networks Trained with the Logistic Loss

URL: View paper

Brief Assessment

Implicit Bias Logistic[58] analyzes two-layer neural networks with logistic loss for classification tasks, not language model pre-training or maximum likelihood estimation for sequence generation. The paper focuses on margin-based generalization in neural networks rather than coverage profiles for autoregressive models.

9. The real-world-weight cross-entropy loss function: Modeling the costs of mislabeling

URL: View paper

Brief Assessment

Real-World-Weight Cross-Entropy[57] focuses on cost-sensitive classification with real-world financial impacts and class imbalance, not on generalization bounds for maximum likelihood estimation or coverage profiles in language model pre-training.

Contribution 3: Algorithmic interventions with provable coverage benefits

Description: The authors propose and analyze three types of interventions: tournament-based model selection procedures that improve upon cross-entropy selection, gradient normalization schemes that achieve horizon-independent coverage bounds, and test-time training strategies that provably enhance coverage for token-level SGD.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Structured convergence through latent epoch reshaping for reordering intermediate computations in large language model training

URL: View paper

Brief Assessment

Latent Epoch Reshaping[38] focuses on gradient accumulation and computational reordering in LLM training, not on model selection procedures, gradient normalization for coverage bounds, or test-time decoding strategies as described in the original paper's contributions.

2. Decoding Large Language Models: An exhaustive guide to understanding, implementing, and optimizing LLMs for NLP applications

URL: View paper

Brief Assessment

Decoding LLMs[40] focuses on general LLM architecture and implementation details (normalization, residual connections) rather than theoretical interventions for coverage optimization in pre-training and post-training contexts.

3. Semantic Image Inpainting with Multi-Stage Feature Reasoning Generative Adversarial Network.

URL: View paper

Brief Assessment

Multi-Stage Feature Reasoning[45] focuses on image inpainting using generative adversarial networks with dynamic partial convolution and feature normalization techniques. This is a computer vision application unrelated to language model training, model selection procedures, gradient normalization for RL, or test-time decoding strategies.

4. Input normalized stochastic gradient descent for language tasks

URL: View paper

Brief Assessment

Input Normalized SGD[37] focuses on gradient normalization for language model training with input normalization techniques, not on tournament-based model selection, test-time training strategies, or coverage bounds for language models.

5. Model Hemorrhage and the Robustness Limits of Large Language Models

URL: View paper

Brief Assessment

Model Hemorrhage[39] focuses on performance degradation during model deployment modifications (quantization, pruning, decoding adjustments), not on pre-training interventions for improving coverage profiles or post-training success metrics.

6. Uni-Perceiver v2: A Generalist Model for Large-Scale Vision and Vision-Language Tasks

URL: View paper

Brief Assessment

Uni-Perceiver v2[44] focuses on multi-modal vision and vision-language task modeling with a unified architecture, not on language model optimization techniques like tournament-based model selection, gradient normalization, or test-time training strategies for coverage improvement.

7. Patch-Growing Universal Adversarial Perturbation

URL: View paper

Brief Assessment

Patch-Growing Perturbation[42] focuses on adversarial perturbation generation for image classification models, not on model selection, gradient normalization, or decoding strategies for language models. The technical domains are entirely distinct.

8. Impartial Multi-task Representation Learning via Variance-invariant Probabilistic Decoding

URL: View paper

Brief Assessment

Variance-Invariant Decoding[43] focuses on multi-task representation learning through probabilistic decoding and variance normalization, not on model selection, gradient normalization, or test-time training strategies for language models as described in the original contribution.

9. Interleaved gradient attenuation through synthetic lexeme currents for structural stability in large language model representations

URL: View paper

Brief Assessment

Interleaved Gradient Attenuation[36] focuses on gradient normalization for structural stability in decoder representations, not on model selection procedures, coverage bounds, or test-time training strategies for language models as described in the original paper.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] The Coverage Principle: How Pre-Training Enables Post-Training View paper
- [1] Limo: Less is more for reasoning View paper
- [2] Kimi k2: Open agentic intelligence View paper
- [3] Will pre-training ever end? a first step toward next-generation foundation mllms via self-improving systematic cognition View paper
- [4] Direct Post-Training Preference Alignment for Multi-Agent Motion Generation Models Using Implicit Feedback from Pre-training Demonstrations View paper
- [5] RobustVLA: Robustness-Aware Reinforcement Post-Training for Vision-Language-Action Models View paper
- [6] GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers View paper
- [7] Large Language Models Post-training: Surveying Techniques from Alignment to Reasoning View paper
- [8] Persistent Pre-Training Poisoning of LLMs View paper
- [9] Mobilellm-r1: Exploring the limits of sub-billion language model reasoners with open training recipes View paper
- [10] Advancing Mathematical Reasoning in Language Models: The Impact of Problem-Solving Data, Data Synthesis Methods, and Training Stages View paper
- [11] Self-improvement in language models: The sharpening mechanism View paper
- [12] LiNeS: Post-training Layer Scaling Prevents Forgetting and Enhances Model Merging View paper
- [13] Predicting post-training reactions from pre-training attitudes View paper
- [14] Image-free Classifier Injection for Zero-Shot Classification View paper
- [15] Post-Deployment Adaptation with Access to Source Data via Federated Learning and Source-Target Remote Gradient Alignment View paper
- [16] A Survey on Post-training of Large Language Models View paper
- [17] Demystifying Domain-adaptive Post-training for Financial LLMs View paper
- [18] Understanding and Improving Length Generalization in Recurrent Models View paper
- [19] A Survey of Post-Training Scaling in Large Language Models View paper
- [20] Poutine: Vision-Language-Trajectory Pre-Training and Reinforcement Learning Post-Training Enable Robust End-to-End Autonomous Driving View paper
- [21] Improving Retrieval-Based Dialogue Systems: Fine-Grained Post-training Prompt Adaptation and Pairwise Optimization Fine-Tuning Strategy View paper
- [22] Conformalized prediction of post-fault voltage trajectories using pre-trained and finetuned attention-driven neural operators. View paper
- [23] A field study of computer efficacy beliefs as an outcome of training: the role of computer playfulness, computer knowledge, and performance during training View paper
- [24] Intelligent optimization of transfer learning and knowledge tracking in enterprise strategic decision support systems View paper
- [25] EvoLM: In Search of Lost Language Model Training Dynamics View paper
- [26] On the Interplay of Pre-Training, Mid-Training, and RL on Reasoning Language Models View paper
- [27] Information Extraction-Driven Knowledge-Based Salient Masked Pre-Training Method View paper
- [28] : Provably Optimal Distributional RL for LLM Post-Training View paper
- [29] Optimising Language Models for Downstream Tasks: A Post-Training Perspective View paper
- [30] Improving the Post-Training Neural Network Quantization by Prepositive Feature Quantization View paper
- [31] Reinforcement Learning with Inverse Rewards for World Model Post-training View paper
- [32] Self-supervised Pre-training for Precipitation Post-processor View paper
- [33] Post Training in Deep Learning with Last Kernel View paper
- [34] Improving Residue-Level Sparsity in RNS-based Neural Network Hardware Accelerators via Regularization View paper
- [35] A Robustly Optimized BERT Pre-training Approach with Post-training View paper
- [36] Interleaved gradient attenuation through synthetic lexeme currents for structural stability in large language model representations View paper
- [37] Input normalized stochastic gradient descent for language tasks View paper
- [38] Structured convergence through latent epoch reshaping for reordering intermediate computations in large language model training View paper
- [39] Model Hemorrhage and the Robustness Limits of Large Language Models View paper
- [40] Decoding Large Language Models: An exhaustive guide to understanding, implementing, and optimizing LLMs for NLP applications View paper
- [41] Straight to the gradient: Learning to use novel tokens for neural text generation View paper
- [42] Patch-Growing Universal Adversarial Perturbation View paper
- [43] Impartial Multi-task Representation Learning via Variance-invariant Probabilistic Decoding View paper
- [44] Uni-Perceiver v2: A Generalist Model for Large-Scale Vision and Vision-Language Tasks View paper

- [45] Semantic Image Inpainting with Multi-Stage Feature Reasoning Generative Adversarial Network. View paper
- [46] TFDP: Token-Efficient Disparity Audits for Autoregressive LLMs via Single-Token Masked Evaluation View paper
- [47] Language modeling techniques for biological sequence processing View paper
- [48] TrajTok: Technical Report for 2025 Waymo Open Sim Agents Challenge View paper
- [49] Human-Inspired Learning for Large Language Models via Obvious Record and Maximum-Entropy Method Discovery View paper
- [50] CoVeR: Conformal Calibration for Versatile and Reliable Autoregressive Next-Token Prediction View paper
- [51] On the Generalization Ability of Next-Token-Prediction Pretraining View paper
- [52] TokenUnify: Scaling Up Autoregressive Pretraining for Neuron Segmentation View paper
- [53] A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses View paper
- [54] On the sample complexity of next-token prediction View paper
- [55] Optimal regularization for a data source View paper
- [56] Beyond maximum-likelihood training: analysis and methods for building robust language generation models View paper
- [57] The real-world-weight cross-entropy loss function: Modeling the costs of mislabeling View paper
- [58] Implicit Bias of Gradient Descent for Wide Two-layer Neural Networks Trained with the Logistic Loss View paper
- [59] Moment distributionally robust tree structured prediction View paper
- [60] Maximizing entropy on adversarial examples can improve generalization View paper
- [61] On how to avoid exacerbating spurious correlations when models are overparameterized View paper
- [62] Limits of sensing temporal concentration changes by single cells View paper